

Strathclyde

Discussion Papers
in Economics



Estimating Public Preferences on Population Health Ethics

Rory Allanson, University of Strathclyde
Matthew Robson, Erasmus University Rotterdam

No. 24 – 2

October 2024

Department of Economics
University of Strathclyde, Glasgow

Estimating Public Preferences on Population Health Ethics

Rory Allanson [†] Matthew Robson ^{‡§}

October 28, 2024

Abstract

We develop a social choice experiment to estimate public preferences on population ethics. Our experiment poses three within-subject treatments in which participants allocate scarce resources to determine the health-related quality-of-life, and existence, of two population groups. Within a flexible social welfare function, we estimate participant-level preferences for inequality aversion, average vs total welfare maximisation, and minimum ‘critical level’ thresholds. By combining random behavioural and random utility models we also explicitly model ‘noise’ in decision making. Using a sample of UK adults (n=115, obs.=5,060), we find that 98.7% of respondents are inequality averse, prioritising the worst-off at the expense of efficiently maximising overall health. The modal group of participants (39.2%) maximise total welfare and have a critical level threshold of zero, however there is extensive heterogeneity in participants’ population preferences. We then demonstrate how these preferences can aid policymaking, where difficult trade-offs emerge between equity and efficiency, average and total welfare, and population size.

Keywords: Experiment, Health, Social Welfare, Inequality, Population Ethics.

JEL Codes: C90, D63, I18.

[†]Corresponding Author. Department of Economics, University of Strathclyde, 199 Cathedral Street, G4 0QU, Glasgow, Scotland. Email: rory.allanson@strath.ac.uk, ORCID: [0000-0001-8759-3598](https://orcid.org/0000-0001-8759-3598). Website: www.roryallanson.weebly.com.

[‡]Erasmus School of Economics and Erasmus Centre for Health Economics Rotterdam, Erasmus University Rotterdam, Burgemeester Oudlaan 50, 3062 PA Rotterdam, The Netherlands; Tinbergen Institute, Gustav Mahlerplein 117, 1082 MS Amsterdam, The Netherlands. Email: robson@ese.eur.nl. ORCID: [0000-0003-4558-7637](https://orcid.org/0000-0003-4558-7637). Website: www.mrobson92.com.

[§]This study was funded by Wellcome Trust (205427/Z/16/Z) and the Erasmus Initiative Smarter Choices for Better Health. Ethics approval was given by the Internal Review Boards of the Erasmus Research Institute of Management and Erasmus School of Economics (2019/10/30-40140tvo and ESE IRB-E Application 2021-10). We would like to thank Stijn Bruers, Richard Cookson, Tim Doran, Olivier L’Haridon, Tom Van Ourti, Owen O’Donnell and Erik Schokkaert, as well as participants at various seminars and conferences, for their helpful comments. We provide an online repository, with estimated preference parameters and code for policy evaluation at: <https://data.mendeley.com/datasets/6wdpxx2hfp>. We have no conflicts of interest to disclose. All errors remain our own.

1 Introduction

Population ethics, a field straddling the boundary between philosophy and economics, grapples with the complex moral and economic questions surrounding issues of welfare, inequality, and population size (Caviola et al., 2022). While authors in the field have considered a range of alternative ‘population principles’ and their consequences for social welfare (Blackorby et al., 2005; Arrhenius, 2012; Parfit, 2016), there remains sparse evidence on the preferences of the public with regards to distributive justice and health. We hence design a novel experiment containing distributional decision problems which allow us to estimate public preferences on population ethics, within a social welfare function framework.

In our experiment, participants assume the role of a social decision maker and must distribute a finite budget of resources to determine the health-related quality-of-life of two hypothetical population groups. Decisions are made across three within-subject treatments. The first two, *fixed-population* treatments, require participants to make distributional decisions across rounds with (in)equally productive and (in)equally populous groups. This forces them to trade-off equity, efficiency, and priority for larger populations. We estimate participant-level preference parameters in order to quantify what trade-offs the public are willing to make. This extends the existing experimental literature, which is limited to fixed-population settings (Dolan and Tsuchiya, 2011; Hurley et al., 2020) where decision makers cannot alter population size (Schokkaert and Tarroux, 2022; Cadham and Prosser, 2023).

In the final *variable-population* treatment, participants can change the size of the total population. In each round one group is initially non-existent, and participants must choose whether or not to bring them into existence when making their resource allocations. This novel existence mechanic allows for the estimation of population preference parameters identifying whether participants prioritise total or average welfare (Blackorby et al., 2005), and whether participants would bring additional persons into existence even if they could not achieve a minimum (critical) level of health (Parfit, 1984). These preferences hence provide empirical evidence of the willingness of the public to sacrifice the welfare of existing populations in favour of prospective persons (Luyten et al., 2022).

The use of experiments in eliciting preference parameters is well-documented. Schokkaert and Tarrow (2022) provide an extensive review of their use in the study of income and welfare, while Costa-Font and Cowell (2019) and Cadham and Prosser (2023) summarise their use in health. As a primary, essential good, health is an important outcome with which to consider public preferences (Anand, 2002), and one which lends itself to experiments better than welfare due to relative ease of measurement (Gaertner and Schokkaert, 2012). Indeed several aspects of health, including its interactions with income and aversion to health inequality, have been well explored experimentally (Ali et al., 2017; Hurley et al., 2020; Robson et al., 2024). A shared feature of these works however is their limitation to fixed-population scenarios.¹ By allowing for a dichotomous choice over existence and thus total population size, our variable-population treatment can relax this constraint.

Beyond the bounds of economics, philosophers have also devoted extensive thought to the questions surrounding health, distributional justice, and population size. While Blackorby et al. (2005) provide a comprehensive overview of these contributions, most relevant to our exercise is those works proposing *population principles*, decision rules which dictate how resources should be deployed to maximise social welfare (Arrhenius, 2012). Of particular interest are three forms of (generalised) utilitarianism - total, average, and critical level - which all provide intuitive alternatives for the prioritisation of resource allocation in society (Parfit, 2016). However as Parfit (1984) addresses, these principles can lead to undesirable outcomes like the ‘Repugnant Conclusion’, which creates an ethical dilemma when considering which is preferable. Experiments can provide us with the public perspective as to which principle aligns best with their preferences, an important contribution as much of the philosophical work in population ethics lacks an empirical basis (Thomson, 2001).² Accordingly, we exploit experimental methods to provide policy-relevant preference parameters towards the principles of population ethics.

¹While authors in other fields have previously utilised incentivised laboratory experiments to investigate how preferences respond to changing group sizes across rounds (Charness and Rabin, 2002; Andreoni, 2007; Fisman et al., 2007; Macro and Weesie, 2016), participants cannot change the groups sizes within rounds.

²Recent authors such as Caviola et al. (2022) and Schönegger and Grodeck (2022) have begun incorporating experiments to answer these questions, but their works do not go as far as to estimate preference parameters.

In order to estimate these parameters, we designed an interactive online experiment populated by a sample ($n=115$) of the UK adult population recruited via Prolific. Across 22 rounds of decision problems and three within-subject treatments, participants distributed resources between two population groups, producing a rich dataset of 5,060 ($115 \times 22 \times 2$) observations. This dataset enables us to estimate participant-level preference parameters within a Social Welfare Function (SWF), provided we first specify a functional form for our SWFs. To this end we select a flexible functional form that stems from Atkinson (1970) and Blackorby et al. (2005), and incorporates parameters for health inequality aversion, ε , alongside our population principles: total vs average, β , and critical level thresholds, τ . Our specified form allows for flexibility in estimation, quantification of difficult trade-offs, and the evaluation of social welfare under alternative policies, all of which are key when contributing policy relevant evidence on the questions of population ethics.³

Importantly, we assume that participants in all rounds make the valuations upon which they base their optimal allocations with some degree of random noise, undoing a common simplifying constraint that participants allocate perfectly in line with their preferences (Robson et al., 2017; Cookson et al., 2018; McNamara et al., 2021). Noise is explicitly modelled with a random behavioural model, which assumes that allocations are drawn from an error-prone distribution unique to each participant, which returns the optimal allocation only on average (Conte and Moffatt, 2014; Robson, 2021). As current works often restrict the use of noise to pooled estimates (Edlin et al., 2012; Hurley et al., 2020), this work represents an extension to the use of noise in experimental health economics. We additionally employ a random utility model (McFadden, 1973; McFadden, 1981) in the latter half of the experiment to account for evaluation error in the existence decision due to noise in the respective welfare valuations. Our estimates hence acknowledge that participants likely make mistakes in both their valuation of allocations and in their subsequent decision making, an additional advancement of the current implementation of estimation noise within the field.

³Note that while the axioms and assumptions underpinning SWFs can be debated theoretically, such theoretical arguments are less convincing for identifying specific parameter values. An experiment like ours is thus important in providing empirical estimates of public preferences.

Our results reveal that a large majority of participants are willing to sacrifice overall population health in order to reduce health inequalities by prioritising the worse-off. We find a median health inequality aversion parameter of 31.10, with the majority (74.68%) of participants classified as *prioritarian* ($1 \leq \varepsilon < 500$); smaller minorities are classified as *efficiency-seeking* (1.27%, $\varepsilon < 1$) or *maximin* (24.05%, $\varepsilon \geq 500$). Regarding population principles, a majority of participants (63.3%) are found to maximise total welfare ($\beta = 1$), compared to 21.5% who alternatively maximise average welfare ($\beta = 0$). Furthermore, 45.6% of participants are estimated to have a critical threshold of zero ($\tau = 0$), whilst the remaining 54.4% have positive thresholds ($\tau > 0$). Importantly however, the distributions of β and τ are found to be non-independent: 94.1% of average welfare maximisers have a positive critical level threshold, compared to only 38.0% of those who maximise total health.

Together, our estimated health inequality aversion and population principle parameters can be used to evaluate a range of real world policies. We demonstrate this with an illustrative example featuring four hypothetical alternative policies, which each affect both the distribution of health and population size. In addition, we provide our estimates and code in an [online repository](#) to enable others to evaluate policies in their own settings.

In sum, we make two main contributions to collective understanding of public preferences towards population ethics. The first, our novel experimental design, extends current work in fixed-population settings (Dolan and Tsuchiya, 2011; Cookson et al., 2018; Hurley et al., 2020; Robson, 2021) to those with variable-populations. This allows for our estimation of public preferences towards varying population size, and analysis of how these preferences compare and interact with those regarding group resource-productivity. Our third treatment then grants the further ability to estimate whether total or average utilitarianism is preferred by participants, and whether participants require it achievable that prospective persons can reach certain critical levels of health before bringing them into existence. These features expand upon mainstream study of health inequality aversion and the equity-efficiency trade-off (Costa-Font and Cowell, 2019; Cadham and Prosser, 2023; Robson et al., 2024). By offering the ability to consider variable-population settings, our design may also prove useful

when studying other preferences that cannot be captured by a fixed-population experiment, such as those in environmental or sociological contexts.

Second, we contribute policy-relevant parameter estimates of public preferences towards the questions of population ethics. While previous authors have estimated public preferences for inequality aversion and the equity-efficiency trade-off (Dolan and Tsuchiya, 2011; Robson et al., 2017; Cookson et al., 2018), our experiment allows for the evaluation of a wider range of important parameters within this field. We therefore take an important step towards the alignment of the philosophical foundations of population health ethics with the empirical evidence provided by experimental economics. As such, the novel experimental evidence our exercise provides can be used to aid difficult policy decisions concerning trade-offs between equity, efficiency and population size.

The rest of this work is structured as follows. Section 2 presents and justifies our experimental design, while Section 3 outlines the theoretical foundations of this work. After Section 4 describes our dataset, Section 5 provides a complete presentation of the results of the experiment. Section 6 then provides discussion, before Section 7 concludes, together considering the relevancy of this work to policymaking and the wider literature. Additional supplementary material provided with this work contains further experimental resources, sensitivity tests, and some supporting information.

2 Experimental Design

We design an experiment to identify the trade-offs the public are willing to make between efficiency, equity, and population size. In the experiment, participants distribute scarce resources between two population groups to determine the health of hypothetical individuals in society. The experiment begins with an interactive tutorial, concludes with a short questionnaire, and consists of three within-subject treatments: Treatments A and B have five scenarios each, while Treatment C has twelve.⁴ The experiment was programmed in R

⁴A short tutorial is also provided before Treatments B and C.

Shiny, and populated using a representative sample of the UK adult population recruited from Prolific. Participants completed the experiment in a single session and were required to answer all questions. While the complete script of the experiment tutorials is provided in Appendix A, this section presents an overview of the experimental setup, and details the specifics of the three treatments.

2.1 Experimental Overview

In each scenario, participants assume the role of a social decision maker (SDM) and must allocate resources to determine the health of hypothetical individuals (i) across two population groups (j). In each group, of population size n_j , all individuals share a common baseline health level y_j and multiplier p_j . Participants then allocate resources x_j , drawn from a budget m , to each group in order to determine the health-related quality-of-life ($q_{ij} = q_j = y_j + p_j x_j$) of the individuals they contain. Health-related quality-of-life is explained to participants using a *Visual Analogue Scale* (VAS) and is bounded between the best (100) and worst (0) health that participants can imagine.⁵ Across 22 scenarios and three within-subject treatments, the budget, multipliers, baseline health, population size, and group existence are exogenously varied, while the length of life for each individual remains constant at 80 years.⁶ Our design thus presents participants with a series of constrained optimisation problems, wherein preferences are inferred from the allocations participants make.


Participants completed the experiment using the user interface shown in Figure 1. The top table provides scenario-specific information about the two groups, including their population size, multiplier, and health-related quality-of-life. Population size is represented by the bar width, and after allocating resources to the groups using the sliders at the bottom, the resulting health-related quality-of-life outcomes are visualised by the bar height. More specifically, the baseline health of a group is represented in light grey, while the additional

⁵We focus on health-related quality-of-life, rather than quality-of-life more generally (Parfit, 1984; Parfit, 2016), as there is a) an extensive experimental literature concerned with eliciting health-related quality-of-life, using Visual Analogue Scales, and b) an extensive applied literature concerned with evaluating alternative health policies using Distributional Cost-Effectiveness Analysis (Cookson et al., 2020).

⁶Participants are told that health-related quality-of-life is an individual’s average across their life.

health they receive from resources is shown in dark grey. The total and remaining budgets are then emboldened on the left, with the total, average, and minimum levels of health across the whole population emboldened on the right. Finally, a minimum timer of 20 seconds per scenario was included to encourage participants to sufficiently consider their choices.

Figure 1: Experiment User Interface



2.2 Treatments

Participants faced 22 scenarios across three within-subject treatments; Treatments A and B have *fixed-population* scenarios, whilst Treatment C has *variable-population* scenarios. Table 1 provides a complete set of example scenarios one participant might see, demonstrating how the multipliers, budget, baseline health, population size, and group existence differ across scenarios and treatments.⁷ Each treatment introduces an additional level of complexity, meaning that randomisation of treatment order across participants was infeasible, and hence each participant completed Treatments A, then B, then C. The order of scenarios within each treatment was, however, randomised. Details of the three treatments are below.

⁷Note that by design, the multipliers are orthogonal to the budget, baseline health, population size, the screen position of a group (left or right), and group existence.

Table 1: Example Experimental Factors

Treatment	Scenario	Multiplier, p_i		Budget	Baseline Health, y_i		Pop. Size (M)		
		Left	Right		Left	Right	Left	Right	Exist
A	1	1	0.50	90	10	10	3	3	Yes
A	2	1	1	60	40	40	3	3	Yes
A	3	0.25	1	30	40	40	3	3	Yes
A	4	0.50	1	30	70	70	3	3	Yes
A	5	1	0.25	90	10	10	3	3	Yes
B	6	1	1	30	10	10	5	3	Yes
B	7	1	0.50	60	40	40	2	4	Yes
B	8	1	0.25	30	40	40	3	1	Yes
B	9	0.25	1	60	10	10	4	5	Yes
B	10	0.50	1	30	70	70	5	1	Yes
C	11	1	1	30	70	0	3	2	No
C	12	0.25	1	90	10	0	2	5	No
C	13	1	0.50	60	40	0	1	4	No
C	14	1	1	90	10	0	2	4	No
C	15	0.50	1	30	70	0	5	3	No
C	16	1	1	30	10	0	3	1	No
C	17	1	0.50	60	10	0	1	2	No
C	18	1	0.25	60	40	0	5	4	No
C	19	0.50	1	90	10	0	4	1	No
C	20	1	1	30	40	0	2	3	No
C	21	1	0.25	30	70	0	1	5	No
C	22	0.25	1	60	10	0	4	2	No
Mean		0.77	0.77	53.18	33.18	15.45	3	2.91	

Notes: *Left* and *Right* refer to the groups on the left and right of the user interface respectively. *Exist* denotes whether both groups initially exist.

2.2.1 Treatment A

In each scenario of Treatment A, the population size and baseline health of the two groups are identical while the budget and multipliers vary. The former varies between scenarios, and the latter between groups and scenarios. In each scenario, group multipliers p_j are pairs drawn randomly from the set $G \in \{(0.25, 1), (0.5, 1), (1, 1), (1, 0.5), (1, 0.25)\}$, with the budget and baseline health of the two groups similarly drawn randomly from the sets $E \in \{30, 60, 90\}$ and $D \in \{10, 40, 70\}$, respectively. These latter two factors are however constrained such that the groups cannot achieve a health-related quality-of-life outcome greater than the

upper limit of 100.⁸ The group population size meanwhile is drawn randomly from the set $T \in \{1, 2, 3, 4, 5\}$; participants are faced with values of T in either exclusively millions or exclusively tens of millions. Overall, these scenarios force participants to make trade-offs between efficiency and equity, as participants must choose between efficiently allocating resources to groups with higher multipliers to maximise total/average health, and reducing health inequality by allocating (some) resources to groups with lower multipliers. The choices participants make hence reveal their degree of health inequality aversion.

2.2.2 Treatment B

Treatment B departs from Treatment A by mandating that the groups be unequally-sized within each scenario.⁹ By posing problems between unequally-sized groups, Treatment B presents more complex allocation decisions which amplify the trade-off between total/average and minimum health in scenarios where larger (smaller) groups also possess a higher (lower) multiplier. We thus employ the results of Treatment B in concert with those of Treatment A to assess participant preferences for inequality aversion, alongside study of participants preferences towards allocating resources to larger or smaller populations.

2.2.3 Treatment C

Treatment C consists of *variable-population* scenarios, where participants can choose to change the population size. In each scenario, one group does not exist initially: participants must, therefore, choose whether or not to bring them into existence and how to subsequently allocate resources.¹⁰ Participants hence have three options: bring the second group into existence and give them resources, bring them into existence but give them nothing,¹¹ or prevent

⁸I.e. a group cannot have a baseline health of 70 and a multiplier of 1 while the budget is 90, as they could then achieve 160 health-related quality-of-life, exceeding the maximum; accordingly, the budget is always 30 when the groups start with 70 baseline health.

⁹Remark that Treatments A and B only have *fixed-population* scenarios as whilst populations sizes may differ, participants cannot choose to vary the population size.

¹⁰Figure A1 illustrates this aspect of the user interface.

¹¹Theoretically undesirable as minimum and average health decrease while total health is unchanged.

their existence and give all resources to the pre-existent group.¹² Across Treatment C there are twelve scenarios, which use a block randomisation design (see Appendix A.4) to ensure that the budget, multipliers, population size and baseline health of the pre-existent group are orthogonal from one another. This forces participants to make trade-offs not only between efficiency and equity, but between average and total welfare, and as to the minimum acceptable level of welfare. In sum, these scenarios enable us to observe the circumstances under which participants will increase population size, and to identify which population principles participants adhere to.

2.3 Experiment Particulars

Prolific, an online sampling platform, was used to recruit and pay participants. The experiment ran as a single session between August 18th-21st 2022, taking 30.34 minutes on average per participant.¹³ Participants were each paid £5 for completing the experiment. In total, the sample consists of 115 participants, whose characteristics are described in Table 2 of Section 4. While 144 began the experiment, 6 did not fully complete it, resulting in a 4.16% attrition rate. Furthermore, 23 participants did not answer at least 3 of the 5 tutorial assessment questions (see Appendix A.2.2) correctly, and hence were excluded from the main sample to help ensure those included sufficiently understood the exercise. To explore if this has any effect on our results, Appendix C.4 conducts sensitivity analysis by providing structural estimates by exclusion criteria.

Before the main experiment, a pilot was conducted between July 25th-August 5th 2022.¹⁴ Consisting of 12 volunteers, the pilot collected feedback used to gauge which areas of the experiment needed refinement. While this feedback was largely positive, several resultant changes were indeed made: notably, Treatment C was assigned 12 rather than 10 questions, and the information on each side of the interface was emboldened.

¹²For each participant the pre-existent group always inhabits one side of the interface: left or right.

¹³Participants took 0.4 minutes longer per year of age to complete the experiment on average, significant at the 1% level; no other demographic variable significantly influenced completion time.

¹⁴More information is provided in Appendix A.3.

3 Theory

3.1 Social Welfare

We assume social decision makers (SDM) aim to maximise social welfare, which is itself assumed to be a function of the weighted sum of transformed health, h_i , of individuals, i , within a population of size N , with population preference parameters β and τ . Accordingly, the SWF takes the following general reduced form:

$$W = N^\beta [h_{EDE} - \tau_h], \quad (1)$$

where

$$h_{EDE} = U^{-1} \left(\sum_i^N \frac{1}{N} U(h_i) \right). \quad (2)$$

The health outcome of individual i , $h_i = q_i l_i$, results from the product of health-related quality-of-life, q_i , and length of life, l_i . The equally distributed equivalent level of health, h_{EDE} ,¹⁵ is determined by the evaluation function of the SDM which transforms each individual's health outcome into a corresponding value of well-being, $U(\cdot)$. β then distinguishes between population principles relating to *average* ($\beta = 0$) or *total* ($\beta = 1$) welfare, while the critical level parameter, τ , determines the minimum desirable health threshold $\tau_h = \tau l$.¹⁶

The evaluation function $U(\cdot)$ allows for alternative parametric functional forms. These functional forms are assumed to be concave, wherein the greater the degree of concavity, the higher the aversion to inequalities in health outcomes. If $U(\cdot)$ is linear (and $\tau = 0$), eq. (1) represents a Utilitarian SWF wherein the SDM evaluates each individual equivalently. Conversely, if $U(\cdot)$ is strictly concave, it represents a Prioritarian SWF which draws diminishing marginal welfare from the health of each successive individual. Taken together, our

¹⁵This variable denotes the level of health such that the social decision maker is indifferent between the current distribution of health and the equal distribution of h_{EDE} across the entire population: $\sum_i^N \omega_i U(h_i) = \sum_i^N \omega_i U(h_{EDE})$.

¹⁶The inclusion of τ ensures that the addition of any individuals for whom $h_i < \tau_h$ contributes negatively to social welfare.

framework for measuring social welfare allows us to capture SDM preferences relating to: a) health inequality aversion, b) total vs. average welfare, and c) critical levels.

In order to estimate preference parameters using our experimental data, two further elements must be specified in this framework. The first is the parametric functional form of $U(\cdot)$, to which we select the iso-elastic functional form assumed in the Atkinson index, namely $U(h_i) = \frac{h_i^{1-\varepsilon} + 1}{1-\varepsilon}$, where $\varepsilon \geq 0, \neq 1$ (Atkinson, 1970).¹⁷ The second is to define the number and nature of the groups $j \in K$ within the population. We specify groups of population size, n_j , within which individuals share identical levels of health, $h_{ij} = h_j, \forall i \in [1 : n_j]$. This simplifies the above SWF to the group level, accounting for the relative size of each group, n_j/N , where h_{EDE} is defined as:

$$h_{EDE} = \left(\sum_j^K \frac{n_j}{N} h_j^{1-\varepsilon} \right)^{\frac{1}{1-\varepsilon}} \quad (3)$$

3.2 Optimal Choices

Experimental participants, acting as the SDM, must exhaust a budget m of resources x_j by distributing them between two groups of individuals such to determine the health outcome of each individual in these groups. Health, which results from the resources afforded to each group, is given in terms of quality-adjusted life expectancy $h_j = q_j l_j = (y_j + x_j p_j) l_j$, wherein q_j is a function of baseline health-related quality of life y_j , the resources allocated to the group x_j , and the health productivity of individuals within the group p_j . The SDM, subject to the binding budget constraint $m = \sum_{j=1}^K x_j$, thus selects the following optimal allocations that result in health levels $h_j^* = (y_j + x_j^* p_j) l_j$ which maximise social welfare W :¹⁸

¹⁷We choose the iso-elastic functional form as Robson et al. (2024) finds it to fit participant's choices better than an exponential functional form (Kolm, 1976).

¹⁸We additionally impose non-negativity conditions on this optimal allocation, to ensure the optimal allocations are feasible within our experimental setup. If $x_j^* < 0$ then we force $x_j^* = 0$ and $x_k^* = m, \forall j \in (1, 2), \neq k$.

$$x_j^* = \frac{m - \frac{y_j l_j}{p_k l_k} \left(\frac{n_k p_k l_k}{n_j p_j l_j} \right)^{\frac{1}{\varepsilon}} + \frac{y_k}{p_k}}{1 + \frac{p_j l_j}{p_k l_k} \left(\frac{n_k p_k l_k}{n_j p_j l_j} \right)^{\frac{1}{\varepsilon}}}, \quad \forall j \in (1, 2), \neq k. \quad (4)$$

Note that for a fixed-population of size N , neither β nor τ affect the optimal allocations. Moreover, the optimal allocations for W are the same as for h_{EDE} . In other words, it is only when a decision will affect population size that β and τ must be accounted for.

In Treatment C, participants are given the choice of whether or not to bring a second group into existence before they allocate resources. If they choose to not bring in the second group, they must exhaust their budget such that $x_j^{**} = m$, and thus $h_1^{**} = (y_1 + mp_1)l_1$; if they do bring in the second group the optimal allocations are as in eq.(4) except with a second optimal allocation $x_j^*, \forall j \in (1, 2)$ required. The choice, D , of whether to bring the second group into existence ($D = 1$) or not ($D = 0$), depends therefore on which delivers the highest welfare, given the optimal allocations above. Participants identify the difference between these welfare levels as V^* :

$$V^* = W_1^* - W_2^* = n_1^\beta (h_1^{**} - \tau_h) - N^\beta \left(\left(\sum_j^K \frac{n_j}{N} h_j^{*1-\varepsilon} \right)^{\frac{1}{1-\varepsilon}} - \tau_h \right). \quad (5)$$

Thus, $D = 0$ if $V > 0$ and $D = 1$ if $V < 0$. In sum, the choice of whether to bring another group into existence depends not only on the optimal health levels - themselves a function of m, y_j, p_j, l_j, n_j , and ε - but also on population preferences β and τ .

3.3 Estimation

The data gathered from Treatments A and B are used to estimate the health inequality aversion parameter, ε . Given $\hat{\varepsilon}$, data from Treatment C are then used to estimate β and τ . We perform this process by combining the optimal allocations from A and B with a random behavioural model, and the choice to bring a group into existence or not from C with a random utility model. Estimating by this procedure allows for the modelling of observed

behaviour as a function of the predicted optimal behaviour, whilst accounting for noise in decision making. Parameters are then estimated using maximum likelihood estimation.

3.3.1 Random Behavioural Model

Observed allocation decisions in Treatment A and B are modelled using a random behavioural model. We assume that the resource share to the first group, $\tilde{x}_1 = x_1/(x_1 + x_2)$, is drawn from a two-limit Tobit distribution $\tilde{X}_1 \sim TOB(\tilde{x}_1^*, \sigma_1)$, following Andreoni and Miller (2002), where \tilde{x}_1^* is the optimal allocation (eq.(4)) and σ_1 is a *noise* parameter. For observations between the limits (0 and 1) the likelihood function for a Tobit draws from a normal probability density function, but for those at the limits, a cumulative density function is employed instead. In our two group allocation problems the resource share allocated to the first group is assumed to be drawn from the Tobit, which ensures all allocations participants could make with non-zero probability are feasible in our estimation.

For each participant, r , we hence find the parameter values that maximise the log-likelihood function defined over the number of rounds $t \in T$ in Treatments A and B:

$$LL_r^{TOB} = \sum_{t=1}^T \left[\sum_{0 < \tilde{x}_{1rt} < 1} \log \left(\frac{1}{\sigma_{1r}} \varphi \left(\frac{\tilde{x}_{1rt} - \tilde{x}_{1rt}^*}{\sigma_{1r}} \right) \right) + \sum_{\tilde{x}_{1rt} \leq 0} \log \left(\Phi \left(\frac{\tilde{x}_{1rt} - \tilde{x}_{1rt}^*}{\sigma_{1r}} \right) \right) + \sum_{\tilde{x}_{1rt} \geq 1} \log \left(1 - \Phi \left(\frac{\tilde{x}_{1rt} - \tilde{x}_{1rt}^*}{\sigma_{1r}} \right) \right) \right], \quad (6)$$

where Φ is the standard normal cumulative distribution function and φ is the standard normal probability density function.

3.3.2 Random Utility Model

In Treatment C, participants first decide whether or not to bring the second group into existence. Making this decision *optimally* will depend on V^* , the difference between the highest social welfare that could be achieved when only the first group exists minus the highest social welfare that could be achieved if both groups existed ($V^* = W_1^* - W_2^*$). If

$V^* > 0$, the SDM would not bring the second group into existence ($D = 0$), but if $V^* < 0$, then they would ($D = 1$). As we assume that participants evaluation of welfare is subject to noise, then the difference between the two options must itself also be noisy. The probability of bringing the second group into existence, $P(D = 1)$, thus depends on the degree of noise, σ_2 , as well as the optimal welfare levels. In sum, the participant is assumed to evaluate the welfare of each group as:¹⁹

$$W_1 = W_1^* + \epsilon_1, \quad W_2 = W_2^* + \epsilon_2. \quad (7)$$

Which gives:

$$\begin{aligned} P(D = 1) &= Pr(W_2 \geq W_1) \\ &= Pr(W_2^* + \epsilon_2 \geq W_1^* + \epsilon_1) \\ &= Pr(\epsilon_1 - \epsilon_2 \leq W_2^* - W_1^*). \end{aligned} \quad (8)$$

We assume that ϵ_1 and ϵ_2 are i.i.d. errors which follow a Gumbel (or Type 1 Generalized Extreme Value) distribution, with mean zero and scale parameter, σ_2 : $\epsilon_1, \epsilon_2 \sim Gumbel(0, \sigma_2)$. Given this assumption, the difference between ϵ_1 and ϵ_2 follows a logistic distribution, with mean zero and scale parameter σ_2 : $\epsilon_1 - \epsilon_2 \sim Logistic(0, \sigma_2)$. This allows us to evaluate a logistic cumulative distribution function, $F(W_2^* - W_1^*; 0, \sigma_2)$, to estimate $P(D = 1)$.

For each participant, r , we then estimate population preferences β_r and τ_r and the scale parameter σ_{2r} using the previously estimated preference parameters, $\hat{\epsilon}_r$, and maximise the following log-likelihood function given the observed choices D_{rt} :

$$LL_r^U = \sum_{t=1}^T \log \left((1 - D_{rt}) \left(1 - \frac{1}{1 + e^{(W_{1rt}^* - W_{2rt}^*)/\sigma_{2r}}} \right) + D_{rt} \left(\frac{1}{1 + e^{(W_{1rt}^* - W_{2rt}^*)/\sigma_{2r}}} \right) \right) \quad (9)$$

3.3.3 Exclusion Criteria

In order to measure the goodness-of-fit of the structural model, after our main estimation we calculate Mean Proportional Likelihood (MPL) values for each participant from LL_r^{TOB}

¹⁹Note that in our estimation we normalise welfare levels by a constant, which does not affect welfare orderings but aids comparability of σ_2 parameters between participants.

and LL_r^U . More specifically, define $PL_t = L_t/(L_t + L_t^{UNI})$, where L_t is the likelihood in round t for the data and estimates, and L_t^{UNI} is a likelihood for a uniform distribution draw. $MPL = 1/T \sum_t (PL_t)$. As $MPL \rightarrow 1$, data fit improves; if $MPL = 0.5$, the fit of the model to the data is no better than fit uniform distribution draws. Accordingly, we hence exclude participants with an $MPL \leq 0.5$ for either the random behavioural or random utility model from our main structural analysis. For more details, and the distribution of MPL for included and excluded participants, see Appendix C.4.

4 Data

Our main analysis draws on the responses of 115 experiment participants. As each participant completed 22 distribution problems posed between two population groups, there are $(115 \times 22 \times 2) = 5,060$ observations for each of the experimental factors described in Section 3. We provide an overview of each of these variables in Table 2.

Table 2: Data Overview

Variable	Notation	Mean	Obs	Range
Budget	m	50.50	5,060	[30-90]
Resources	x_j	25.25	5,060	[0-90]
Resource Share	\tilde{x}_i	0.50	5,060	[0-1]
Multiplier	p_j	0.77	5,060	[0.25-1]
Relative Multiplier	\tilde{p}_j	0.50	5,060	[0.20-0.80]
Health-Related Quality-of-Life	q_j	40.97	5,060	[0-100]
Health Share	\tilde{q}_j	0.50	5,060	[0-1]
Baseline Health	y_j	21.89	5,060	[0-70]
Population Size	ω_j	17.38	5,060	[1-50]
Relative Population Size	$\tilde{\omega}_j$	0.50	5,060	[0.17-0.83]
Both Exist	-	0.76	5,060	[0-1]

Note: *Both Exist* captures whether both groups were given existence within a scenario (trivially true in Treatments A and B).


In addition to this experimental data, each participant also answered a short questionnaire that enquired about some ancillary issues and requested simple demographic information; as participants could choose to not answer each questionnaire item, the number of responses varies slightly. We detail the results of this questionnaire in Appendix B.

5 Results

5.1 Fixed-Population Treatments

Using data from Treatment A, Figure 2 plots the distributions of resource shares (\tilde{x}_j) and health shares (\tilde{q}_j) conditional on the relative multiplier (\tilde{p}_j). This illustrates the equity-efficiency trade-offs that participants are willing to make in the experiment.²⁰

Figure 2: Resource and Health Share Distributions, by Relative Multiplier



Note: Empirical cumulative density function over all participants ($n = 115$) and rounds in Treatment A (obs. = 1150) of *resource shares* ($\tilde{x}_j = x_j / \sum x_j$) and *health shares* ($\tilde{q}_j = q_j / \sum q_j$) to groups distinguished by relative multipliers ($\tilde{p}_j = p_j / \sum p_j$).

As efficiency mandates that participants allocate resources to the more productive (i.e. higher multiplier) group, the extent to which participants prioritise efficiency can be seen through their (lack of) prioritisation for more productive groups. As shown by the left panel, strong evidence is found that participants allocate against efficiency: in only 50% of rounds did participants allocate 20%+ of resources to groups four times as productive as their scenario counterparts ($\tilde{p}_i = 0.80$). Conversely, groups that were four times less productive

²⁰Where $\tilde{p}_j = \frac{p_j}{p_j + p_k}$. As multipliers are drawn in pairs as either (0.25, 1), (0.50, 1), (1, 1), (1, 0.50), or (1, 0.25), \tilde{p}_j is either 0.20, 0.33, 0.50, 0.67, or 0.80 respectively.

($\tilde{p}_i = 0.20$) were assigned 80%+ of resources in more than half of such rounds. The majority of participants hence allocated resources to prioritise the less productive, and hence reduced inequalities in health rather than efficiently maximising total health.

The equalisation of health outcomes is further demonstrated by the right panel, which presents proportional health shares by relative multiplier. Observe the highly prominent jumps at 0.5, the point where health is equalised between the groups: most participants strived for equitable health distributions irrespective of relative group productivity. More productive groups were however assigned a health share above 0.5 by a significant minority of participants (top right), which also makes intuitive sense: although these groups were assigned significantly less resources than their counterparts on average, they use those resources more efficiently. This demonstrates that in terms of health outcomes there was at least some trade-off with distributional efficiency.

Sample averages for the influence of multipliers and population size towards allocations can also be generated using linear regression; these estimates are contained in Table 3. Regressions (1) and (3) confirm the results of Figure 2: groups with larger relative multipliers are allocated significantly fewer resources by participants on average. This result is consistent with equation (4), which for inequality-averse respondents ($\varepsilon > 1$) predicts a decrease in x_{jA}^* when p_j (and hence \tilde{p}_j) increases.²¹ Regressions (2) and (4) then confirm that more health is assigned to groups with larger multipliers, though with smaller coefficients than those of regressions (1) and (3). This follows similarly from eq.(4), which predicts a small (large) increase in q_{jA}^* for inequality-averse (efficiency-seeking) participants when p_j increases.

Our Treatment B results also detail that larger population groups are awarded both more resources and health than their scenario counterparts on average. As population size functions as a second, indirect productivity factor in our experimental framework, this result is logical as devoting resources to larger groups serves the efficient maximisation of both total and average welfare. The estimated coefficient for relative multipliers in regression (3) is then negative, as in (1). It thus remains true in Treatment B that the majority of participants

²¹Inversely, participants who prioritise efficiency ('efficiency-seeking', $\varepsilon < 1$) are predicted to increase allocations to more productive groups.

Table 3: Resource and Health Shares by Relative Multiplier and Population Size

	Treatment A		Treatment B	
	(1) Resources Coef./ <i>(S.E.)</i>	(2) Health Coef./ <i>(S.E.)</i>	(3) Resources Coef./ <i>(S.E.)</i>	(4) Health Coef./ <i>(S.E.)</i>
Relative Multiplier	-0.646*** (0.06)	0.138*** (0.02)	-0.524*** (0.06)	0.171*** (0.02)
Relative Population Size			0.327*** (0.05)	0.141*** (0.02)
Constant	0.500*** (0.00)	0.500*** (0.00)	0.336*** (0.02)	0.430*** (0.01)
N	115	115	115	115
Observations	1150	1150	1150	1150
Overall R^2	0.459	0.177	0.335	0.269

Notes: * = $p < 0.10$; ** = $p < 0.05$; *** = $p < 0.01$. Relative multipliers are re-centred around 0.50 to improve coefficient interpretation. A random-effects model with robust standard errors was employed to counteract within-participant error clustering. Demographic control variables are omitted from these regressions due to negligible coefficients and for presentation purposes; see A.5 for discussion.

prioritise equity by reducing the gap in health outcomes between the two groups, despite often doing so in defiance of relative group productivity.

5.2 Variable-Population Treatment


5.2.1 Existence Decisions

Participants chose to bring the second group into existence in 769 of 1,380 Treatment C scenarios (55.27%). Their willingness to do so can be best understood by considering how often each participant brought them in out of the 12 rounds, as presented by Figure 3.

Only 6 participants, 5.22% of the sample, never brought the second group into existence, while 19 participants took the inverse position to always bring in the second group.²² The remaining 78.26% of participants brought the second group into existence in only some rounds. To understand what motivated these choices, in Appendix C.1 we utilise a random

²²The block randomisation design employed in Treatment C ensured that all participants faced scenarios where bringing the second group into existence was efficient, both in terms of average and total health. This choice to never bring a second group into existence is, therefore, likely a moral or heuristic position. This was also suggested by pilot feedback (see A.3.)

Figure 3: Participant-Level Existence Decisions



Notes: Distribution of participant-specific existence decisions, averaged across the 12 rounds of Treatment C for each participant. Reveals the proportion of rounds each participant brought the second group into existence. Distribution is shown as both a histogram, with density normalised to 1, and an empirical cumulative density plot.

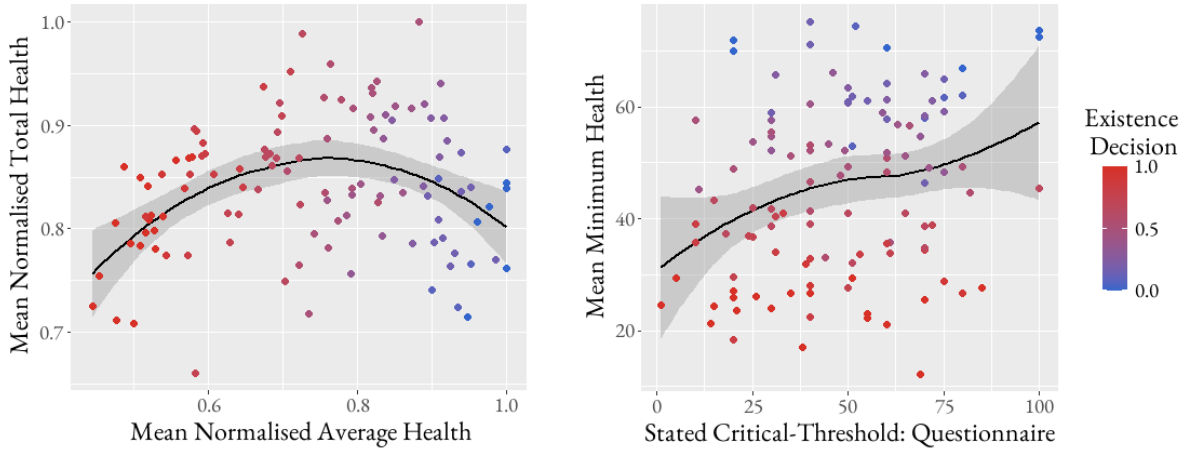
effects model to explore what experimental variables affected this decision most. Consistent with our population principles, we find that relative multipliers and relative population sizes each play an important part in how participants approach this decision, with larger and, particularly, more resource-productive groups brought in frequently. An increasingly generous budget, finally, had a positive but comparatively minor effect on the willingness of our participants to grant existence to the second group.

5.2.2 Population Principles

In Treatment C, the decision to bring the second group into existence depends on the trade-offs that participants are willing to make between maximising total, average, and minimum welfare. This is a separately interesting question from the equity-efficiency trade-off in Section 5.1.²³ To begin, Figure 4 plots participant-level mean total, average and minimum levels of health, alongside participant’s average existence decisions (of Figure 3).

²³As proved by Blackorby et al. (2005), in fixed-population treatments like A and B the maximisation of total welfare and average welfare is equivalent.

Figure 4: Existence Decisions with Average, Total and Minimum Health Trade-offs



Notes: Participant markers are coloured by the proportion of times they brought the second group into existence: blue if never (0), red if always (1). Fitted fractional polynomial line is shown in black, with associated 95% Confidence Interval in grey.

To better gauge what portion of respondents sought to maximise either average or total health, the left panel of Figure 4 plots the two against one another. Total and average health are normalised at the round level, and then averaged across all rounds for each participant; a value of 1 therefore means that a participant made allocations in each round so to maximise either total or average health, respectively. The panel illustrates that there is a clear trade-off between the two principles, wherein the existence decision determines which can be maximised. More specifically, participants who never bring the second group into existence (blue) are better able to maximise average health. In contrast, those who brought them in only sometimes (maroon) are generally better at maximising total health, but in doing so must sacrifice some average health. Participants who always bring the second group into existence (red) meanwhile effectively maximise neither total nor average health.

The right panel of Figure 4 then displays how the existence decisions of participants affected mean minimum health levels across the experiment, and the relation of these levels to their stated critical-thresholds.²⁴ Our first observation is that participants who brought the

²⁴These are elicited from participant’s responses to the question “Imagine you could choose whether an individual would come into existence or not [...] what is the minimum level of health that individual would have to have in order for you to bring them into existence?”.

second group into existence less often (blue) are better able to maximise minimum health across experimental rounds. This is logical insofar as they chose to do so when the prospective population was significantly more resource productive than their scenario counterparts, making it highly efficient to bring them in. Second, participants' stated critical-thresholds are positively correlated with the minimum level of health that participants awarded on average, and negatively correlated with their decisions to bring in the second group. For more information and analysis on the stated critical-thresholds, see Appendix C.2.


5.3 Preference Parameter Estimates

While our non-structural results outline the heterogeneity with regards to the population principles that participants exhibit in their choices, the proper disentanglement of which principles participants preferred - whilst accounting for their inequality aversion - requires the structural estimation of these parameters at the participant-level. Accordingly, we derive estimates of health inequality aversion from the allocation problems of Treatments A and B, ε , while our population preference parameters - average-total, β , and critical-thresholds, τ - are estimated using existence decisions from Treatment C. All structural results are shown for our analytical structural sample ($n = 79$); further details on exclusion criteria and structural estimates for excluded participants are found in Appendix C.4.

5.3.1 Health Inequality Aversion

Figure 5 details the distribution of participant-level health inequality aversion parameters, ε . Overall, participants exhibit substantial aversion to health inequalities, with a median ε of 31.10. There is however extensive heterogeneity in the degree of aversion: few participants (1.27%) are classified as Efficiency Seeking ($\varepsilon < 1$), while the majority (74.68%) are classified as Prioritarian ($1 \leq \varepsilon < 500$); a minority (24.05%) are classified as Maximin ($\varepsilon \geq 500$). The vast majority of participants are, therefore, substantially averse to health inequalities.

Figure 5: Distribution of Health Inequality Aversion, ε




Notes: Distribution of participant-specific ($n = 79$) health inequality aversion parameters, ε , estimated using data from Treatment A and B. Distribution is shown as both a histogram, with density normalised to 1, and an empirical cumulative density plot. The x-axis is shown on a log-scale, ε values over 500 are censored.

5.3.2 Population Parameters

Next, Figure 6 plots the participant-level distribution of average-total, β , and critical threshold, τ , population preference parameters. Regarding average-total preferences, results show that the majority of participants (63.3%) maximise *total* welfare ($\beta = 1$), whilst 21.5% instead maximise *average* welfare ($\beta = 0$). The remainder, 15.2%, maximise a mixture of average and total welfare ($0 < \beta < 1$). For critical thresholds, the majority of participants (54.4%) have a positive critical-threshold ($\tau > 0$), while a sizeable minority (45.6%) have a critical threshold of zero ($\tau = 0$). Of those with $\tau > 0$, the mean critical threshold is 19.41.

An interesting question is then to enquire whether the distributions of β and τ are independent. Cross-tabulating these parameters, Table 4 reveals that they are not: of those participants who maximise average welfare ($\beta = 0$), the vast majority (94.1%) have positive critical thresholds ($\tau > 0$). Within the majority that instead maximised total welfare, we find that only 38.0% have a positive critical-threshold. In sum, our findings show that while

Figure 6: Distribution of Population Parameters, β and τ



Note: Distribution of participant-specific ($n = 79$) average-total parameters, β , and critical threshold parameters, τ estimated using data from Treatment C. Distribution is shown as both a histogram, with density normalised to 1, and an empirical cumulative density plot.

the modal group of participants maximise total welfare, with no critical threshold (39.2%), there is substantial heterogeneity in public preferences towards population ethics.

Table 4: Population Preferences Cross-Tabulation

		Average-Total, β		
		$\beta = 0$	$0 < \beta < 1$	$\beta = 1$
		<i>21.5%</i>	<i>15.2%</i>	<i>63.3%</i>
Threshold	$\tau = 0$	<i>45.6%</i>	1.3%	5.1%
	$\tau > 0$	<i>54.4%</i>	20.3%	10.1%

Note: Percentage of participants ($n = 79$) classified by average-total, β , and critical threshold, τ , parameter estimates. Presented with italics for tabulated and normal font for cross-tabulated categorisations.

5.4 Policy Choice

The public preference parameters we estimate above have direct relevance towards evaluating health and demographic policy. To evidence this claim, Table 5 presents a simple

illustrative example. Consider four alternative policies - A, B, C, and D - which impact the health-related quality of life and existence of four prospective persons; selecting between these policies involves difficult trade-offs between equity and efficiency, average and total health, minimum levels of health, and population size. However by employing our estimated preference parameters, we can predict the level of welfare that each participant in our experiment would assign to each policy. By then using the respective social welfare valuations W_r (eq. 1) made by each participant, we can identify the percentage of our sample who would ‘vote’ for each policy as their preferred option.

Table 5: Alternative Policies and their Associated Outcomes

Policy	Persons				Summary		
	1	2	3	4	Average	Total	Gap
A	20	100			60	120	80
B	40	60			50	100	20
C	40	60	50		50	150	20
D	40	60	50	10	40	160	50

Note: Valuations are in terms of health-related quality-of-life, in the same fashion as the main experiment.

We predict that of 79 participants, 0%, 7.6%, 92.4%, and 0% would vote for Policies A, B, C, and D respectively as their most preferred option. These results follow from the mean welfare values across all participants, equal to 16.31, 43.56, 62.68, and 4.68 for our four respective policies. We draw three immediate conclusions from these findings. Firstly, from the general preference for Policy B over Policy A, we understand that in a fixed-population setting participants are willing to sacrifice average (or total) health in order to improve the health of the worst-off in society. Secondly, the relative popularity of C over B leads us to conclude that in a variable-population setting, our participants prefer the inclusion of an additional person provided that this person would have a health level *at least as good as* the average health of the existing society. Finally from our result that no individual in our sample would vote for policy D - despite this option having the highest total health - we make the important note that for the majority of participants with $\beta = 1$, is is not total

health, but total *welfare* which participants strive to maximise. Given high aversion to health inequality and/or positive critical-thresholds, the inclusion of an additional individual with very low health hence actually reduces social welfare.²⁵

Simple illustrations such as this one can be surprisingly reflective of the real world. For example, three prominent areas of contemporary policymaking, namely public health, family planning, and climate change, all have the potential to affect both population size and the distribution of health within a given population. By utilising the preferences parameters we estimate in Section 5, policymakers may evaluate the consequences for social welfare stemming from any policy for which the effects on the distribution of health are known. In an [online repository](#), we provide code for a “policy evaluation function”, which allows others to use our estimated preference parameters to evaluate alternative health policies of their choosing. We do remark however that these evaluations are not a prescriptive statement of what ought to be done. Instead, the descriptive empirical evidence we provide should be considered as an informative tool for debate and policy choice when difficult trade-offs emerge, one which allows for public preferences to be acknowledged in decision making.

6 Discussion

At the fundamental level, we argue that there are at least three reasons why we should value public preferences as an input towards distributive decisions. Firstly, in democratic societies where individuals contribute towards institutions like the healthcare system, we find it natural that their preferences for how resources are deployed within those systems be a point of consideration of policymakers. While expert opinion remains essential in designing large institutions, we believe that public opinion can contribute by outlining guiding principles for how such systems should operate. Our second argument follows that in fields featuring

²⁵This framework allows us to also test the ‘repugnant conclusion’ of Parfit (1984) by considering one final policy, Policy E, which is identical Policy D except for the inclusion of X additional persons with a health-related quality-of-life of 1. We can then ask: is there a finite number X for which Policy E would be preferred to any of the other policies? The answer is no for 53.16% of our 79 participants. We hence conclude that the majority do indeed find the ‘repugnant conclusion’, repugnant.

morally complex decisions that affect health or welfare - where there is often no objectively best solution to be found - the integration of the preferences of those who will be subject to said decisions is integral towards reaching fair and equitable outcomes. Otherwise, the pursuit of efficiency may overshadow what the public would prefer as an ideal distribution (Luyten et al., 2022). Finally, the parameters we provide here can assist philosophers in reaching reflective equilibrium as to which population principles align most closely with true public preferences. Insofar as works such as ours depend on such principles in their design, providing philosophers with a robust empirical basis of parameter estimates can assist in their creation of the theoretical platform upon which we build.

In our exercise, we find substantial heterogeneity between participants as to their preferred equity-efficiency trade-off and population principle. As illustrated by Figures 2 and 5, the majority of respondents prefer equitable distributions over efficient ones, in that resources are often allocated to relatively inefficient health producing groups. This conclusion is consistent with the wider literature. Specifically, Robson et al. (2017) find that 81.51% of English adults would sacrifice efficiency to reduce health inequalities²⁶ using an online survey, while Edlin et al. (2012), Attema et al. (2015), and Ali et al. (2017) find similar inequality-averse preferences in UK adults using various alternative methods. Participants were then similarly divided on the total vs average welfare trade-off in Treatment C: Figure 6 shows that a minority (21.5%) prefer to maximise average welfare at the expense of total, while a large majority (63.3%) took the reverse position. This divide echoes that found previously by Caviola et al. (2022), although there is scope for further research in this area.²⁷

Our results also lend credence to the concerns of Luyten et al. (2022) that health policy guidelines - often constructed to maximise total welfare - may not fully capture public preferences, as evidence of a contrasting alignment with average welfarism is found in Figure 6. In concert with the supporting evidence of Section 5.1 on health inequality aversion, we conclude that health policies that prioritise efficient deployment of healthcare resources could better serve public desires by, at least partly, promoting equity in health outcomes.

²⁶See also Dolan and Tsuchiya (2011) and Abásolo and Tsuchiya (2013).

²⁷For example into other population principles such as number-dampened utilitarianism (see Ng, 1986).

The finding of considerable willingness by participants to bring in prospective populations via the existence mechanic of Treatment C is also policy relevant, as it raises the question as to if policymakers sufficiently consider the welfare of future or otherwise prospective populations (Neumann et al., 2016). The actualist and comparativist perspectives of welfare theory (Bykvist, 1998), which incorporate the utility of future persons into valuations of social welfare, may thus be more appropriate than the presentist perspective that does not (Arrhenius, 2005). The choice between these approaches is important as it affects cost-effectiveness evaluations of policies that influence population size (Blackorby et al., 2005), or that incur future non-medical costs (see de Vries et al., 2019). As for the Repugnant Conclusion (Parfit, 1984; 2016), as average welfarism avoids it while total welfarism does not, mixed evidence of public aversion towards it is found.

7 Conclusion

Population ethics presents important questions for the way that policymakers evaluate the welfare of populations both present and future. In order to provide evidence of public opinion towards these questions, we conduct a distributional social choice experiment using a sample of 115 British adults. The experiment challenges participants to solve distribution problems posed between two population groups who can be (un)equally sized, (un)equally productive, and when one of the groups does not initially exist. In doing so, we evaluate how participants approach the equity-efficiency trade-off, whether their allocations align closest with total, average, or critical level generalised utilitarianism, and whether they are willing to sacrifice the welfare of a pre-existent group for that of a prospective one.

We find that 98.7% of participants are willing to sacrifice distributional efficiency by dis-favouring more resource productive population groups in their allocations, instead preferring to promote equity in resulting health outcomes. This conclusion aligns with much of the existing literature (see McNamara et al., 2020), and suggests that most UK adults hold a weighted proiritarian perspective (median ε of 31.10). A majority (63.3%) of participants

and then found to maximise total welfare over average welfare, with less than half (45.6%) having a critical threshold of 0. Participants were most divided in their preferences as to whether or not to bring in the prospective populations of Treatment C, with larger and more productive groups more likely to be brought in.

The empirical evidence presented in this work can contribute to effective policymaking by informing decision makers how those affected by population health decisions would themselves prefer welfare to be allocated. Moreover, the novel existence mechanic introduced in this work furthers existing understanding of preferences for the welfare of persons future as well as present. We believe that these contributions could prove fruitful for authors across the range of disciplines that engage with population ethics, and encourage the possibility of further, multidisciplinary research into the important questions raised herein.

Bibliography

- Abásolo, I. and A. Tsuchiya (2013). “Is more health always better for society? Exploring public preferences that violate monotonicity”. In: *Theory and Decision* 74(4), pp. 539–563. DOI: <https://doi.org/10.1007/s11238-011-9292-1>.
- Ali, S., A. Tsuchiya, M. Asaria, and R. Cookson (2017). “How Robust Are Value Judgments of Health Inequality Aversion? Testing for Framing and Cognitive effects”. In: *Medical Decision Making* 37(6), pp. 635–646. DOI: <https://doi.org/10.1177/0272989X17700842>.
- Anand, S. (2002). “The concern for equity in health”. In: *Journal of Epidemiology and Community Health* 56(7), p. 485. DOI: <https://doi.org/10.1136/jech.56.7.485>.
- Andreoni, J. (2007). “Giving gifts to groups: How altruism depends on the number of recipients”. In: *Journal of Public Economics* 91(9), pp. 1731–1749. DOI: <https://doi.org/10.1016/j.jpubeco.2007.06.002>.
- Andreoni, J. and J. Miller (2002). “Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism”. In: *Econometrica* 70(2), pp. 737–753. DOI: <https://doi.org/10.1111/1468-0262.00302>.
- Arrhenius, G. (2005). “Superiority in Value”. In: *Recent Work on Intrinsic Value*. Berlin: Springer, pp. 291–304. DOI: [10.1007/s11098-004-5223-0](https://doi.org/10.1007/s11098-004-5223-0).
- Arrhenius, G. (2012). “The Impossibility of a Satisfactory Population Ethics”. In: *Descriptive and Normative Approaches to Human Behavior*. World Scientific, pp. 1–26. DOI: https://doi.org/10.1142/9789814368018_0001.
- Atkinson, A. B. (1970). “On the measurement of inequality”. In: *Journal of Economic Theory* 2(3), pp. 244–263. DOI: [https://doi.org/10.1016/0022-0531\(70\)90039-6](https://doi.org/10.1016/0022-0531(70)90039-6).
- Attema, A., W. Brouwer, O. l’Haridon, and J. L. Pinto (2015). “Estimating sign-dependent societal preferences for quality of life”. In: *Journal of Health Economics* 43, pp. 229–243. DOI: <https://doi.org/10.1016/j.jhealeco.2015.07.006>.

- Blackorby, C., W. Bossert, and D. Donaldson (2005). *Population Issues in Social Choice Theory, Welfare Economics, and Ethics*. 39. Cambridge: Cambridge University Press. DOI: <https://doi.org/10.1017/CCOL0521825512>.
- Bykvist, K. (1998). “Changing preferences: A study in preferentialism”. PhD thesis. Uppsala: Uppsala University Publications. ISBN: 99-2853350-4.
- Cadham, C. J. and L. A. Prosser (2023). “Eliciting Trade-Offs Between Equity and Efficiency: A Methodological Scoping Review”. In: *Value in Health* 26(1), pp. 943–952. DOI: <https://doi.org/10.1016/j.jval.2023.02.006>.
- Caviola, L., D. Althaus, A. L. Mogensen, and G. P. Goodwin (2022). “Population ethical intuitions”. In: *Cognition* 218(104941). DOI: <https://doi.org/10.1016/j.cognition.2021.104941>.
- Charness, G. and M. Rabin (2002). “Understanding Social Preferences with Simple Tests”. In: *The Quarterly Journal of Economics* 117(3), pp. 817–869. DOI: <https://doi.org/10.1162/003355302760193904>.
- Conte, A. and P. G. Moffatt (2014). “The econometric modelling of social preferences”. In: *Theory and Decision* 76, pp. 119–145. DOI: <https://doi.org/10.1007/s11238-012-9309-4>.
- Cookson, R., S. Ali, A. Tsuchiya, and M. Asaria (2018). “E-learning and health inequality aversion: A questionnaire experiment”. In: *Health Economics* 27(11), pp. 1754–1771. DOI: <https://doi.org/10.1002/hec.3799>.
- Cookson, R., S. Griffin, O. F. Norheim, and A. J. Culyer (2020). *Distributional Cost-Effectiveness Analysis: Quantifying Health Equity Impacts and Trade-Offs*. Oxford University Press. DOI: <https://doi.org/10.1093/med/9780198838197.001.0001>.
- Costa-Font, J. and F. Cowell (2019). “Incorporating Inequality Aversion in Health-Care Priority Setting”. In: *Social Justice Research* 32, pp. 172–185. DOI: <https://doi.org/10.1007/s11211-019-00328-6>.

- de Vries, L. M., P. H. van Baal, and W. B. Brouwer (2019). “Future Costs in Cost-Effectiveness Analyses: Past, Present, Future”. In: *Pharmacoeconomics* 37(2), pp. 119–130. DOI: <https://doi.org/10.1007/s40273-018-0749-8>.
- Dolan, P. and A. Tsuchiya (2011). “Determining the parameters in a social welfare function using stated preference data: an application to health”. In: *Applied Economics* 43(18), pp. 2241–2250. DOI: <https://doi.org/10.1080/00036840903166244>.
- Edlin, R., A. Tsuchiya, and P. Dolan (2012). “Public preferences for responsibility versus public preferences for reducing inequalities”. In: *Health Economics* 21(12), pp. 1416–1426. DOI: <https://doi.org/10.1002/hec.1799>.
- Fisman, R., S. Kariv, and D. Markovits (2007). “Individual Preferences for Giving”. In: *American Economic Review* 97(5), pp. 1858–1876. DOI: <https://doi.org/10.1257/aer.97.5.1858>.
- Gaertner, W. and E. Schokkaert (2012). *Empirical Social Choice: Questionnaire-Experimental Studies on Distributive Justice*. Cambridge: Cambridge University Press. DOI: <https://doi.org/10.1017/CB09781139012867>.
- Hurley, J., E. Mentzakis, and M. Walli-Attaei (2020). “Inequality Aversion in Income, Health, and Income-related Health”. In: *Journal of Health Economics* 70 (1), p. 102276. DOI: <https://doi.org/10.1016/j.jhealeco.2019.102276>.
- Kolm, S.-C. (1976). “Unequal inequalities. I”. In: *Journal of Economic Theory* 12(3), pp. 416–442. DOI: [https://doi.org/10.1016/0022-0531\(76\)90037-5](https://doi.org/10.1016/0022-0531(76)90037-5).
- Luyten, J., E. Verbeke, and E. Schokkaert (2022). “To be or not to be: Future lives in economic evaluation”. In: *Health Economics* 31(1), pp. 258–265. DOI: <https://doi.org/10.1002/hec.4454>.
- Macro, D. and J. Weesie (2016). “Inequalities between Others Do Matter: Evidence from Multiplayer Dictator Games”. In: *Games* 7(2), p. 11. DOI: <https://doi.org/10.3390/g7020011>.

- McFadden, D. (1973). “Conditional Logit Analysis of Qualitative Choice Behavior”. In: *Frontiers in Econometrics*. Ed. by P. Zarembka. Academic Press: New York, pp. 105–142.
- McFadden, D. (1981). “Econometric Models of Probabilistic Choice”. In: *Structural Analysis of Discrete Data with Econometric Applications*. Ed. by C. Manski and D. McFadden. MIT Press: Cambridge, pp. 198–272.
- McNamara, S., J. Holmes, A. K. Stevely, and A. Tsuchiya (2020). “How averse are the UK general public to inequalities in health between socioeconomic groups? A systematic review”. In: *The European Journal of Health Economics* 21(2), pp. 275–285. DOI: <https://doi.org/10.1007/s10198-019-01126-2>.
- McNamara, S., A. Tsuchiya, and J. Holmes (2021). “Does the UK-public’s aversion to inequalities in health differ by group-labelling and health-gain type? A choice-experiment”. In: *Social Science & Medicine* 269, p. 113573. DOI: <https://doi.org/10.1016/j.socscimed.2020.113573>.
- Neumann, P. J., G. D. Sanders, L. B. Russell, J. E. Siegel, and T. G. Ganiats (2016). *Cost-Effectiveness in Health and Medicine*. Oxford: Oxford University Press. DOI: <https://doi.org/10.1093/acprof:oso/9780190492939.001.0001>.
- Ng, Y.-K. (1986). “Social criteria for evaluating population change: an alternative to the Blackorby-Donaldson criterion”. In: *Journal of Public Economics* 29(3), pp. 375–381. DOI: [https://doi.org/10.1016/0047-2727\(86\)90036-8](https://doi.org/10.1016/0047-2727(86)90036-8).
- Office for National Statistics (2021). *Population estimates for the UK, England and Wales, Scotland and Northern Ireland, provisional: mid-2020*. Titchfield, Hampshire. URL: <https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/populationestimates/bulletins/annualmidyearpopulationestimates/mid2020>.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press. DOI: <https://doi.org/10.1093/019824908X.001.0001>.
- Parfit, D. (2016). “Can We Avoid the Repugnant Conclusion?” In: *Theoria* 82(2), pp. 110–127. DOI: <https://doi.org/10.1111/theo.12097>.

- Robson, M. (2021). “Inequality Aversion, Self-Interest and Social Connectedness”. In: *Journal of Economic Behavior and Organization* 183, pp. 744–772. DOI: <https://doi.org/10.1016/j.jebo.2020.12.029>.
- Robson, M., M. Asaria, R. Cookson, A. Tsuchiya, and S. Ali (2017). “Eliciting the Level of Health Inequality Aversion in England”. In: *Health Economics* 26(10), pp. 1328–1334. DOI: <https://doi.org/10.1002/hec.3430>.
- Robson, M., O. O’Donnell, and T. Van Ourti (2024). “Aversion to health inequality—Pure, income-related and income-caused”. In: *Journal of Health Economics* 94, p. 102856. DOI: <https://doi.org/10.1016/j.jhealeco.2024.102856>.
- Schokkaert, E. and B. Tarrow (2022). “Empirical Research on Ethical Preferences: How Popular is Prioritarianism?” In: *Prioritarianism in Practice*. Ed. by M. Adler and O. Norheim. Cambridge University Press, pp. 459–517. DOI: <https://doi.org/10.1017/9781108691734.010>.
- Schönegger, P. and B. Grodeck (2022). “Concrete Over Abstract: Experimental Evidence of Reflective Equilibrium in Population Ethics”. In: *Issues in Experimental Moral Philosophy*. London: Routledge. URL: <https://ssrn.com/abstract=4128205>.
- Thomson, W. (2001). “On the axiomatic method and its recent applications to game theory and resource allocation”. In: *Social Choice and Welfare* 18(2), pp. 327–386. DOI: <https://doi.org/10.1007/S003550100106>.

APPENDICES

Appendix A Experiment Details

Within the experiment, all participants are shown the below on-screen text for the instructions, tutorial, and tutorial questions. The instructions explain the concept of health-related quality-of-life and provide an overview of the experiment. The seven stages of the tutorial then explain how to use the on-screen interface, and each of the scripts are followed by an interactive on-screen tutorial. Presented next are five tutorial questions which confirm and reinforce participant’s understanding, immediately preceding the three treatments of the main experiment. Two single-slide tutorials are also shown in the transition between treatments, in order to instruct the participant how the problems would change. Lastly, the experiment concluded with a short socio-demographic questionnaire. The full script of these elements is provided here such to provide clarity as to what participants were shown; note that each paragraph that follows contains the information of one slide, with slide titles (where included) highlighted in bold.

A.1 Instructions

Health. Throughout this experiment you will be asked to think about the Health-Related Quality of Life of hypothetical people in society. We will call this “Health”. The scale to the right hand side shows “Health”. The scale is numbered from 0 to 100. 100 means the BEST health you can imagine. 0 means the WORST health you can imagine. Assume that all hypothetical people in this experiment live for exactly 80 years. The “Health” of a person is their average level of “Health” across their life. Please click Next to continue.

Please Read These Instructions Carefully. You will be asked to make decisions which determine the Health of hypothetical people in society. You will be given a “Budget” that you must divide between groups of people. The Budget is the total amount of “Resources”

available to spend. Resources improve Health. Giving more Resources to a group increases the Health of each person within that group. The impact of Resources on Health is determined by a number referred to as the “Multiplier”. The higher the Multiplier, the higher the level of Health each person gains from a given number of Resources. On the screen, you will distribute Resources between two groups of people. You will do this a number of times. Each screen will show a different scenario. The choices you make on one screen will not affect the scenarios that follow. There are no right or wrong answers. We are interested in the choices you make, whatever they are. You will now go through a tutorial, which will explain how to use the computer interface and the exact nature of the experiment. Please click Next to continue.

A.2 Tutorial

A.2.1 Main Tutorial

Tutorial. This tutorial will show you how to use the on-screen interface. You will first be shown the information for one “Group” of people. Groups are identified by letters (e.g. Group S) and have a “Population” of a certain size, in millions of people (e.g. 5M People). Each group begins with an initial level of “Health”. The people within each Group have the same “Health”, but “Health” can vary between Groups. The Group and Population size are shown in a table at the top of the screen. A grey bar in the middle of the screen shows the Population size and initial Health.

The width of the grey bar indicates the Population size of each group. This is also shown as a label at the bottom of the bar. The height of the grey bar, and the grey number to the left of the bar, shows the initial level of Health of each person in that Group. This is measured on a scale numbered from 0 to 100. Where 100 is the BEST health you can imagine, and 0 is the WORST health you can imagine. Please click Next to see this information. If you need to reread this screen press Back. When you are done press Next again.

You will first get practice in giving Resources to only one group of people. Drag the horizontal slider at the bottom of the next screen to the right to give more Resources to people within this group. The amount of Resources you give is shown by the number in the table at the top of the page. Once you have dragged the slider, you can use the left and right arrow keys to make precise changes to the amount. Press the arrow key for a change of 0.1 and hold the arrow key for changes of 1. The Resources you give to a group are taken from the Budget, which is shown on the left of the screen. As you increase the Resources, the Remaining Budget will decrease. You must always use all of the Budget, so that the Remaining Budget is zero. When there is only one Group, this means dragging the slider all the way to the right. Later you will have to distribute the Budget between Groups. Press Next to try out the slider. When you are done, allocate all the Budget and press Next.

The Resources you give to a Group, in addition to their initial Health, determines the “Health” of each person within the group. Health is equal to initial Health, plus the Resources multiplied by a number we call the “Multiplier”. The Multiplier is shown in the table at the top of the screen. Initial Health is shown by the height of the grey bar and the number to the left of the bar. When you give Resources by moving the slider, the resulting Health is shown by the combined height of the grey and black bars. The number to the right of this bar is the amount of Health, and this is also shown in the table at the top of the screen. For example, imagine a group of people who each have an initial Health of 40 and a Multiplier of 1. If you give all of a Budget of 60 resources to that Group the Health of the people within that group will be 100. They will live in the BEST Health you can imagine. Press Next and see how Health changes as you adjust the Resources given to the Group. When you are done, allocate all of the Budget and press Next.

The Multipliers can vary from Group to Group. In the previous scenario, the Multiplier was 1. In the next scenario, it is 0.5. Press Next and see how Health of people with the group changes as you give them more Resources. Notice that there is now a difference between the Resources given and the additional Health achieved. When you are done, allocate all of the Budget and press Next.

In each round of the experiment, there are two groups of people. Groups are identified by letters (e.g. Group B and Group K) and their Population size (in millions, M). On the next screen, there are two sliders at the bottom of the screen that you can use to give Resources to each Group and so determine the Health of the people in those groups. The Health of each person in a Group is shown by combined height of the grey and black bars, where the initial Health is grey and allocated Health in black. The table at the top also shows the Resources, Multiplier and Health for each Group. Now you must allocate the Budget across the two Groups. In doing so, you determine the Health of the people in each group. In the example on the next screen, the Multipliers are the same both Groups. You must use all of the Budget, so that the Remaining Budget (on the left) equals zero. Remember, if you have used the whole Budget, you will not be able to move any slider to the right. If you want to give more Resources to one Group, you will need to give less to another Group first. Press Next and then give Resources to the two Groups. When you have used all of the Budget on the next screen, press Next.

The two Groups change from round to round. On the previous screen, both Groups had a Multiplier of 1. But the Multipliers can differ between Groups, as on the next screen. Move the sliders to give Resources to the two Groups and notice how the Health achieved depends on the Multiplier of each Group. If you give both Groups the same Resources, their Health will differ. Take note of the size of the Budget, which can change from screen to screen. If you are having difficulty seeing both the table and the graph on your screen, zoom out on your web browser by holding “Ctrl” and pressing “-”. Hold “Ctrl” and press “+” to zoom in. Press Next and then give Resources to the two Groups. When you the Remaining Budget is zero, press Next.

The right of the screen shows further summary information. “Total Health” is the total amount of Health across both groups. Intuitively, this is represented by the combined area of the grey and black bars of both groups. “Average Health” is average level of Health per person across the whole population (i.e. $\text{Total Health} / \text{Total Population Size}$). “Minimum Health” is the minimum amount of Health for persons within either Group. This is shown by

the height of the lowest grey/black bar. Press Next and then distribute Resources between the two Groups. Try different allocations and notice how the Total Health, Average Health and Minimum Health depend on your choice. When you are finished, allocate all of the Budget and press Next.

A.2.2 Tutorial Questions

The five assessment questions included at the end of the main tutorial as a competency test are listed below. Possible responses are listed with the correct answer(s) highlighted in bold. After completing the five questions, participants were informed which questions they had answered (in)correctly, and if they were incorrect, why this was.

1. On each screen, you will give Resources to how many Groups of people? - Options: **2**, 3, 4, Not Sure.
2. You can make and adjust the Resources you give by (tick all that apply): - Options: **Clicking and Dragging the Sliders, Using the Arrow Keys**, Moving the Vertical Bar, Not Sure.
3. If you give 60 Resources to a Group with an initial Health of 40 and a Multiplier of 1, what will the Health of people in that Group be? - Options: 40, 70, **100**, Not Sure.
4. If you give 60 Resources to a Group with an initial Health of 40 and a Multiplier of 0.5, what will the Health of people in that Group be? - Options: 40, **70**, 100, Not Sure.
5. Once you have finished giving the Resources, you proceed to the next screen by: - Options: Clicking Next, **Ensuring the Remaining Budget = 0, then Clicking Next**, Waiting, Not Sure.

A.2.3 Treatment B Tutorial

Experiment. You have completed the first part of the experiment. Now you will begin the second part. In the second part the Population Size of the two groups will vary in each

scenario. The Population Size of each group is shown in the table at the top of the screen, by the width of each bar and by the label at the bottom of the bars. You will be given one practice round to explore how differing Population Sizes change the consequences of your allocations. Try allocating different amounts of Resources to the different sized Groups to see these consequences. Press Next to Begin.

A.2.4 Treatment C Tutorial

Experiment. You have completed the second part of the experiment. Now you will begin the third part. In third part, each scenario will start with only one Group of people in existence. You can choose to give all Resources to that Group OR you can choose to bring a second Group of people into existence, and then allocate Resources between the two Groups. You can bring this second Group into existence by ticking the button “Group X Exists?”. They will start with an initial Health level of 0. You can then allocate Resources between the two Groups. You can untick the “Group X Exists?” to not bring that Group into existence. As before, the Population Size and Multipliers will change between these Groups. You will be given one practice round to explore the consequences of bringing another Group of people into existence or not, and of your chosen allocations between the Groups. Try allocating all Resources to the first Group. Then try bringing the second Group into existence and allocating resources between the two Groups to see these consequences. Press Next to Begin.

A.3 Pilot Study

We ran a small pilot study in July 2022 to test and receive feedback on our experimental design, recruiting a convenient sample of twelve participants. While pilot feedback was generally positive, two main modifications were made to the experiment in light of participant’s feedback. First, the number of Treatment C scenarios was increased from ten to twelve. This decision was made primarily to implement a block randomisation design (see A.4), but also because participants did not report fatigue after the experiment. Second, summary statistics

on the right of the user interface were emboldened to promote the visibility of this information. Other modifications were more minor: new options were added to the employment and education questionnaire items, and some wording changes were made in the tutorial.

A.4 Treatment C Block Randomisation

Table A1 show shows the block randomisation design of Treatment C. This design ensures that all participants saw sufficiently distinct combinations of the three random factors, namely population size, multipliers, and baseline health values for the pre-existent group.

Table A1: Treatment C Block Randomisation Design


Multiplier	Large Population		Small Population	
	Low y_i	High y_i	Low y_i	High y_i
(1,1)	1	2	3	4
(1, p_i)	5	6	7	8
(p_i , 1)	9	10	11	12

Remark that each block describes the characteristics of the pre-existent group in that round. Each participant would see one of the twelve types of scenarios (numbered from one to twelve) described in Table A1; they would face these scenarios in a random order decided by a shared seed. As such, there would be six rounds where the pre-existent group would have a ‘Large’ population size, randomly drawing from the set $\{3, 4, 5\}$, and six with a ‘Small’ population size drawn from the set $\{1, 2\}$. Within each of these blocks there would be three rounds where they would have a ‘Low’ baseline health value of 10, and three rounds where they would have a ‘High’ baseline health value drawn from the set $\{40, 70\}$. The final block, multipliers, created three groups of four scenarios: in the first four the groups had an equal multiplier (1,1), in the second four the pre-existent group had a higher relative multiplier (1, p_i), and in the final four they had a lower relative multiplier (p_i , 1), where p_i was drawn randomly from the set $\{0.25, 0.50\}$.

A.5 Existence Interface

Treatment C sees our participants make decisions wherein one of the two groups is initially non-existent. Participants then have the option to bring them into existence, as illustrated by Figure A1. Note that if the box was to be left unchecked, Group K would not exist and their slider would be hidden.

Figure A1: Treatment C Existence Interface



Appendix B Descriptive Statistics

Table B1 details the demographic characteristics of our experimental sample. There is substantial heterogeneity in participants socio-demographic characteristics, with a near 50/50 gender split, an average age close to that of the UK at large (Office for National Statistics, 2021), and a range of educational backgrounds and labour market outcomes.

Table B1: Demographic Data Overview

Variable	Mean	Obs	Range
Female	0.49	115	[0-1]
Married	0.51	115	[0-1]
Age	43.13	109	[19-71]
Income (£K)	29.362	102	[2.5-175]
<i>Highest Education</i>			
Postgraduate	0.23	115	[0-1]
Undergraduate	0.37	115	[0-1]
A-Level	0.29	115	[0-1]
Secondary/Primary	0.11	115	[0-1]
<i>Labour Market Status</i>			
Employed	0.63	115	[0-1]
Unemployed	0.08	115	[0-1]
Retired	0.15	115	[0-1]
Student	0.07	115	[0-1]
Other	0.07	115	[0-1]

Appendix C Additional Results

C.1 Existence Decision

In Table C1, we estimate what factors were most contributive towards participant’s decisions as to whether or not to bring in the initially non-existent groups of Treatment C. Note that in this table the dependant variable ‘Give Existence’ equals one when the second group is given existence and zero otherwise.

Table C1: Factor Influence on Existence Decision

	Give Existence Coef./ (S.E.)
Budget	0.0026*** (0.001)
Relative Multiplier	0.5983*** (0.064)
Relative Population Size	0.2307*** (0.054)
Constant	0.5332*** (0.084)
N	115
Observations	1380
Overall R^2	0.1224

Notes: * = $p < 0.10$; ** = $p < 0.05$; *** = $p < 0.01$. relative multipliers are re-centred around 0.50 to improve coefficient interpretation. A random-effects model with robust standard errors was employed to counteract within-participant error clustering. Demographic control variables are omitted.

Our first takeaway from these results is that initially non-existent groups with a larger multiplier than their pre-existent counterparts were much more likely to be brought in. This finding is consistent with our population principles: those who seek to maximise total welfare will bring them in as they are more resource productive for a given population size, while those who favour equity will do so in order to both balance the health outcomes of the two groups *and* to increase the minimum level of health experienced by either population. Similarly, larger groups were more likely to be brought into existence as they are more productive for a given multiplier, demonstrating that the extension to consider unequally-sized groups has an appreciable impact on how participants allocate resources. Finally, giving

participants a larger budget to award also increased their willingness to grant existence to the second group, with the impact of this factor found to be very close to that of relative population size for rounds with the largest possible budget (90). As such, we understand that in budget-constrained scenarios such as ours, both the characteristics of the two groups and the size of the budget can affect the decision making of participants.

C.2 Stated Critical-Thresholds: Questionnaire


As part of the questionnaire, we elicited *stated* critical-threshold values from participants. These provide us with an alternative to our *estimated* critical-threshold, τ , allowing for a comparison of results. This item presented participants with a visual analogue scale, and requested them to select a response on an interactive slider to the following question:

“Imagine you could choose whether an individual would come into existence or not. If this individual existed, they would live for 80 years. Imagine their Health, averaged across their life, could be captured on a scale numbered from 0 to 100. Where 100 means the BEST health you can imagine, and 0 means the WORST health you can imagine. What is the MINIMUM LEVEL OF HEALTH that individual would have to have in order for you to bring them into existence?”.

Figure C1 plots the resulting distribution of stated threshold values. We find a mean threshold value is 48.08, with only a few participants (4.34%) selecting a threshold below 10. Interestingly, these values are larger on average than our estimated critical threshold parameter τ from Section 5.3.2.

To further explore why this was the case, we next provide linear regression results in Table C2 of the stated critical-thresholds on each participants’ mean existence decision, minimum-level of health in Treatment C, and our estimated critical-thresholds, τ . We observe from these results that the stated thresholds are significantly associated with all three regressors, and in the expected directions: participants with a higher stated threshold are less likely to

Figure C1: Distribution of Stated Critical-Thresholds



Notes: Distribution of elicited stated critical-threshold values, from questionnaire. Distribution is shown as both a histogram, with density normalised to 1, and an empirical cumulative density plot.

bring the second group into existence, more likely to ensure there is a higher minimum level of health across experimental rounds, and have a higher estimated τ . As an additional check, model (4) drop our sample down to 79 participants by applying our exclusion criteria, and finds that while the estimate for τ remains positive, its magnitude and significance decline.

Table C2: Stated Critical-Threshold Regressions

	(1) Stated Threshold Coef./ <i>(S.E.)</i>	(2) Stated Threshold Coef./ <i>(S.E.)</i>	(3) Stated Threshold Coef./ <i>(S.E.)</i>	(4) Stated Threshold Coef./ <i>(S.E.)</i>
Existence Decision	-23.128*** (6.85)			
Mean Minimum Health		0.384*** (0.14)		
Critical Threshold, τ			0.231*** (0.09)	0.128 (0.09)
Constant	60.966*** (4.18)	30.679*** (6.66)	46.402*** (2.25)	51.355*** (2.79)
N	115	115	115	79
R^2	0.106	0.071	0.032	0.014

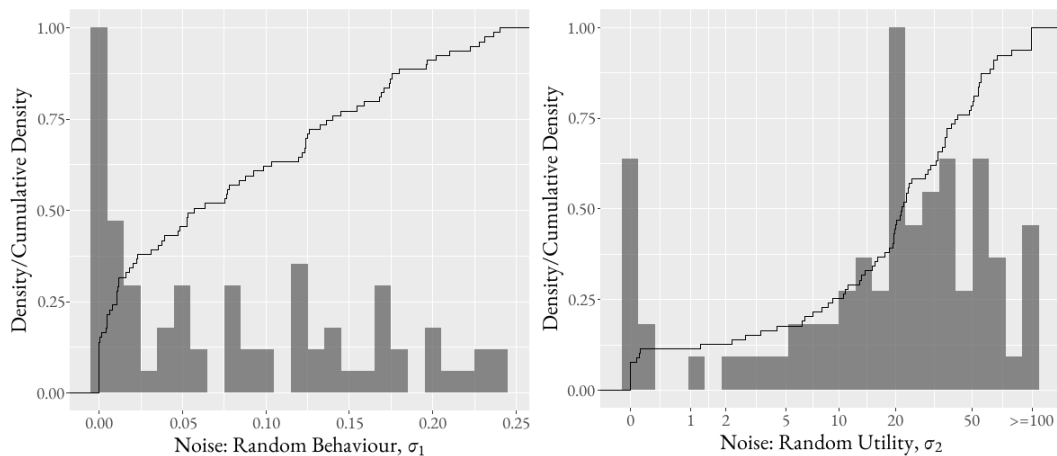
In sum, these results provide assurance that the choices participants made in the experiment are in line with their stated critical-threshold values. However, we do observe a

discrepancy between the estimated and stated critical-threshold values. These differences illustrate the difficulty involved in pinning down precise critical threshold values.

C.3 Noise Parameter Estimates

Figure C2 details the distributions of our two estimated noise parameters, namely σ_1 from the random behavioural model, and σ_2 from the random utility model. Overall, we find a median value for σ_1 of 0.077, and the median value for σ_2 is 38.087. The distributions make it clear however that there is substantial heterogeneity in each estimated parameter.

Figure C2: Distribution of Noise Parameters, σ_1 and σ_2



Notes: Distribution of participant-specific ($n = 79$) noise parameters σ_1 and σ_2 , for the random behavioural and random utility models, respectively. σ_1 is estimated using data from Treatment A and B. σ_2 is estimated using data from Treatment C. Distribution is shown as both a histogram, with density normalised to 1, and an empirical cumulative density plot. The x-axis is shown on a log-scale, σ_2 values over 100 are censored.


C.4 Sample Sensitivity

In our structural analysis, we apply exclusion criteria to improve the quality of the data. More specifically, participants were excluded if they a) did not answer at least 3 of the 5 tutorial questions correctly ($n = 23$), and b) had a MPL lower than 0.5 from Treatment A/B

($n = 6$) or Treatment C ($n = 30$). This left an analytic structural sample of $n = 79$. Here, we present the distribution of estimated preference and noise parameters, alongside mean proportional likelihood values, for participants who were included (blue) and excluded (red).

Figure C3 shows these distributions for our health inequality aversion (ε), average-total (β) and critical threshold (τ) parameters. They detail that included participants were found to have generally higher values for ε and, in particular, a higher proportion of these participants were classified as *maximin* ($\varepsilon \geq 500$). Included participants were also more likely to focus on average welfare ($\beta = 0$) and less likely to have a critical threshold of zero ($\tau = 0$).

Figure C3: Distribution of Preferences Parameters: Included and Excluded




Notes: Distribution of participant-specific ($n = 138$) mean proportional likelihoods, for those those included in (blue) and excluded from (red) our main structural analysis. Distributions are shown as empirical cumulative density plots.

The top panel of Figure C4 plots the distribution of our estimated noise parameters, namely σ_1 from the random behavioural model, and σ_2 from the random utility model. These parameters are found to be generally lower for included participants, as we see a large proportion of excluded participants with high noise parameters ($\sigma_2 \geq 100$).

The bottom panels of Figure C4 plots the distributions of participant-specific Mean Proportional Likelihood (*MPL*) values, which provide an intuitive measure of goodness-of-fit of the structural model. Define $PL_t = L_t / (L_t + L_t^{UNI})$, where L_t is the likelihood in round t for the data and estimates, and L_t^{UNI} is a likelihood for a uniform distribution draw, and $MPL = 1/T \sum_t (PL_t)$. If $MPL = 0.5$, model fit to data is no better than fit uniform

Figure C4: Noise Parameters and Mean Proportional Likelihoods: Included and Excluded



Notes: Distribution of participant-specific ($n = 138$) noise parameters, σ_1 and σ_2 , and mean proportional likelihood values for those included in (blue) and excluded from (red) our main structural analysis. Distributions are shown as empirical cumulative density plots.

distribution draws. As $MPL \rightarrow 1$, data fit improves. We plot the distribution of MPLs in Treatment A and B, in the left panel, and for Treatment C, in the right. Results show that MPLs are, generally, higher for included participants, and no participants were included in the main sample if the structural model fit was as bad or worse than a uniform draw.

Taken together, these figures demonstrate that whilst there are differences in the distributions of estimated preference parameters between the included and excluded participants, we observe lower noise and better goodness-of-fit for included participants. This underlines the importance of our exclusion criteria towards estimating accurately.